



Disaster Recovery with Tiger Bridge

Step-by-step Guide for
Setting up Cloud-enabled
Disaster Recovery
with Tiger Bridge

Version: 1.0
January 2022




Table of Contents

Introduction	3
Why Tiger Bridge?	4
Installation.....	5
Replication.....	10
Recovery	11

[Send Us Your Feedback](#)

[About Tiger Technology](#)

Introduction

Disaster Recovery (DR) comprises the methods that an organization has put in place to respond to and recover from an event, such as a natural disaster or a ransomware attack, which negatively impacts their business operations.

With data at the heart of any organization, DR planning is a vital process which must ensure that the business impact of such events is minimized by:

1. using a method that offers the greatest cost-efficiency, does not require excessive amounts of storage resources, and introduces the lowest amount of IT complexity;
2. ensuring that the two key metrics that measure the effectiveness of the DR method, the Recovery Point Objective (RPO) and Recovery Time Objective (RTO) are kept as short as possible.

From a workflow perspective, a common goal of a DR method will be to enable mission-critical operations, systems, and applications to resume as fast as possible and to ensure access to organizational data based on priority needs.

Read about [Disaster Recovery as a Service \(DRaaS\) in our blog](#).

Why Tiger Bridge?

Tiger Bridge is a software-only data management solution which enables continuous file replication to cloud tiers thus ensuring mission-critical operations can resume immediately in the event of a disaster. Data is stored in a non-proprietary format which means third-party solutions are not required to access it and it can be recovered from the cloud even when Tiger Bridge is not running.

Tiger Bridge replicates new data as it is created, so the Recovery Point Objective (RPO) is as short as the transfer time of the last modified file. This continuous replication removes the need for regular full backups, reducing transmission times and storage costs. The Recovery Time Objective (RTO) is short and recovery from any server failure is easy with Tiger Bridge. In case of partial data loss, users simply access cloud data on demand and resume operations. In case of a complete failure, metadata can be restored from the cloud to a virgin server in minutes and you can restore the data that you need now, which allows you to resume operations with minimum delay. There is no need to wait for the entire backup to be restored to continue work. Data can be fully or partially restored to any number of new or repaired servers. In all cases, data is replicated and synchronized across sources and targets.

Tiger Bridge offers a number of additional benefits such as file server extension and disaster recovery mechanisms, files versioning, cloud migration capabilities, remote collaboration, multi-site synchronization and data protection. Check out our [YouTube channel](#) to learn more.


In this article we will cover the following important steps in the configuration of the Tiger Bridge DR workflow:

1. [Installation](#)
2. [Replication](#)
3. [Recovery](#)


Installation

For all the rest of the steps you will need Tiger Bridge set up and running. In this section, we will outline the needed steps:


1. Open the Bridge installer and follow the first few steps. There is an installation wizard which guides you through the steps.




2. Once this is done, a **Tiger Bridge Configuration** icon will appear on your desktop. Click to open it. You will be able to establish the communication between Tiger Bridge and a cloud provider there. Prior cloud account existence is required. If prompted for a Tiger Bridge license, enter it so you can continue.
3. Click the **Local source** button to specify a folder or drive to be replicated to the cloud.




4. Select the desired drive or folder and click the **OK** button to confirm.



5. Then select the **Add target** button at the bottom to specify a target on the cloud for your data.



6. You can pick any of the available public or on-premises cloud options.



7. Fill in the fields with accurate credentials, such as **Account name** and **Account key**. These may vary depending on the selected provider. Such sensitive information is available to cloud administrators with exclusive rights.

Tiger Bridge target: 

Targets: Azure

Name: Azure

Account name: bridgevideos

Account key:

Blob endpoint: https://bridgevideos.blob.core.windows.net

☒ Use secure transfer (SSL/TLS)

Default access tier: Hot

Rehydration priority: Standard

Containers: List containers

spaces-nas


The available options are summarized in the following table:

Field	Description
Targets	The selected cloud provider
Name	Name of the target – the way it will appear in the Tiger Bridge Configuration
Account name	Your cloud provider's assigned account name
Account key	The secret key for authorization in front of the cloud provider
Use secure transfer (SSL/TLS)	Checked by default for secure communication
Default access tier	The options names vary from one cloud provider to another but you will find the option to select hot, cool, or archive as your default tier. Archive is the cheapest option but retrieving the data is much more expensive and time consuming than the other options.
Rehydration priority	Will be standard by default, but you can change that for some critical sources and give them priority over other standard ones
Buckets/Containers	You will have to select an individual bucket or container on the cloud to sync your data with

Click the **Apply** button once you are ready with your selections.

At the end of this initial configuration Tiger Bridge will be in a paused state waiting for you to finish setting everything up.

8. Next, set your **replication policy**, meaning how frequently the data is backed up. By default, the value is **1 minute**, but it can be days or weeks as well. A new file put within the source directory or changes made to the file/folder structure in general would trigger a replication after the specified value.



9. Select the Settings option on the left and make the needed choices for your environment.

Tiger Bridge settings:

Retrieve mode:

Copy

☒ Delete replica when source file is removed

☒ Keep replica versions

☒ Preserve security descriptor on sync

☒ Allow remote shell extension

☐ Allow partial updates

☒ Enable progressive retrieval

☐ Enable ransomware protection

Maximum modified files to trigger:

0

Event viewer logs:

☐ Replicate file
☐ Move to archive


☐ Reclaim space
☐ Retrieve file

In the following table you can find a brief summary of these settings and their function.


Setting	Function
Retrieve mode	To let Tiger Bridge keep the replica on the target, when the file is successfully retrieved on the source volume, select Copy. To let Tiger Bridge remove the replica from the target, when the file is successfully retrieved on the source volume, select Move.
Delete replica when source file is removed	Select the check box to let Tiger Bridge remove the replica from the target, upon deleting the file from the source volume.
Keep replica versions	By default, if a replicated file is modified on the source and is queued for replication again, the new copy on the target overwrites the previous one. As long as your target supports versioning, you can enable this feature and Tiger Bridge will not overwrite the replica on the target but allow the versioning software to keep each copy as a separate version.
Preserve security descriptor on sync	Select the check box to keep the security of all files on each source after contents synchronization.
Allow remote shell extension	You can install the Tiger Bridge shell extension as a standalone component on a remote computer and thus allow users with administrative privileges to perform manual Tiger Bridge operations on a NAS source or on a local source exported as an SMB share. To be able to use the shell extension from a remote computer, a Tiger Bridge administrator must enable this remote shell extension access setting. Important: Disabling the remote shell extension access only prevents users from executing the commands in the Tiger Bridge context menu in Windows Explorer, but does not restrict their access to data on your source
Allow partial updates	By default, when an already replicated file is modified on the source and needs to be replicated again, Tiger Bridge replicates the whole file on the target. By enabling partial file updates, you let Tiger Bridge overwrite on the target just the parts of the file, which have been modified, if your target supports this functionality.
Enable progressive retrieval	You can disable the progressive retrieval of data and this way configure Tiger Bridge to retrieve only the portion of the file, which is currently being read by the application as long as the respective application supports reading only portions of a file.
Enable ransomware protection	As ransomware attacks usually result in encryption of as many files as possible, Tiger Bridge lets you specify the maximum number of already replicated files, queued for replication, because they have been modified on the source. When the number of modified files queued for re-replication reaches the threshold you have specified, Tiger Bridge automatically pauses all its operations and logs this in the Windows Event Viewer. You can then delete the encrypted files (if there are any), retrieve from the target their unencrypted copies, and then resume normal operations.
Event viewer logs	Select the check boxes for the Tiger Bridge actions you would like to get logged in Event Viewer.

Replication

If Tiger Bridge is still paused after the initial configuration and you visit your cloud provider's web portal, the contents of the selected bucket or container will be **empty**. Open the Tiger Bridge Configuration, click on Tiger Bridge from the menu, and then select **Resume** to initiate replication of your files.



As a result, orange **tags** will appear on all files and folders kept within your source folder, indicating successful replication (or copying) to the cloud.



Replicated contents will also appear on the cloud provider platform, stored in a **non-proprietary** format. They can be downloaded directly without the help of Tiger Bridge.

To see the most up-to-date information with regards to how many files are replicated, in your source folder or drive, select the **Refresh all** button in the Tiger Bridge Configuration and then click on your source.

Replicated:	3 files (244.39 KB)
Nearline:	56 files (1.92 GB)
Archived:	2 files (162.93 KB)
Total managed:	59 files (1.92 GB)



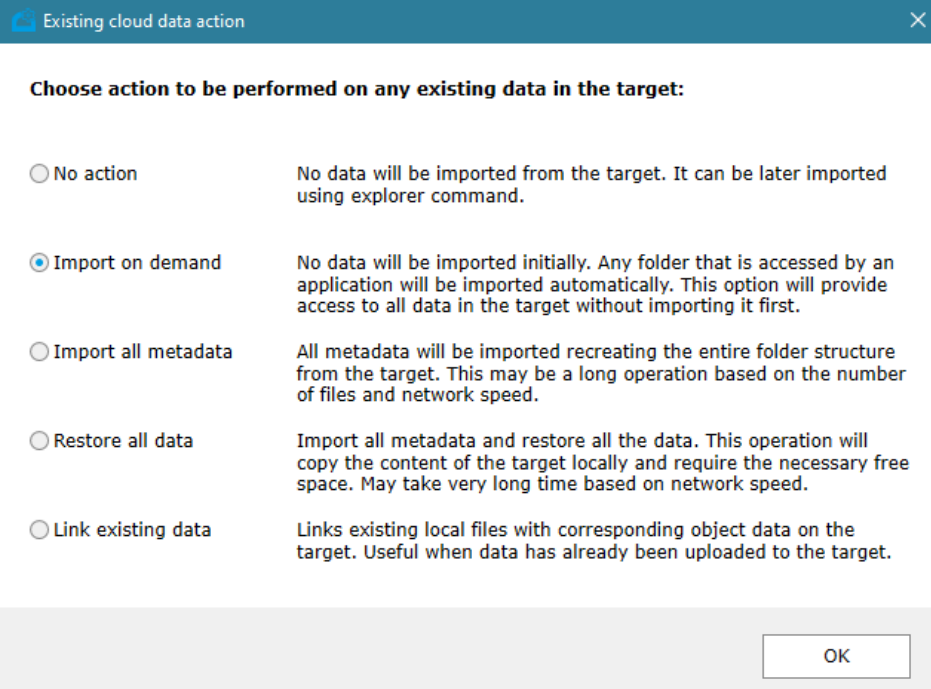
Recovery

In case of an emergency, attack, or a failure, a properly set-up Bridge configuration is a must. Otherwise, if the data is not properly backed up in the cloud, **data loss** becomes inevitable in the case of a disaster.

Data loss is permanent when there are no existing copies of your files elsewhere. However, if your source folder contents were replicated beforehand, they will remain accessible on the cloud, and it is still possible to recover them.

Even if your computer cannot be started, you can install Bridge on another device to get your files back as per the [Installation](#) section of this article:

1. Set it up and activate Tiger Bridge on a new device.
2. Add a **new** source or change it to your previous (and now empty) source directory.
3. Click **OK** and set the target to be the same bucket or container that you used before.
4. Select the appropriate action to be performed on the data that is present on the target:



Existing cloud data action

Choose action to be performed on any existing data in the target:

☐ No action	No data will be imported from the target. It can be later imported using explorer command.
☒ Import on demand	No data will be imported initially. Any folder that is accessed by an application will be imported automatically. This option will provide access to all data in the target without importing it first.
☐ Import all metadata	All metadata will be imported recreating the entire folder structure from the target. This may be a long operation based on the number of files and network speed.
☐ Restore all data	Import all metadata and restore all the data. This operation will copy the content of the target locally and require the necessary free space. May take very long time based on network speed.
☐ Link existing data	Links existing local files with corresponding object data on the target. Useful when data has already been uploaded to the target.

OK

The available options are listed in the following table:

Action	Function
No action	No data will be imported on your source. You can import it later on, if needed.
Import on demand	Tiger Bridge creates a stub file (nearline or offline file, depending on the tier of your target the replica is stored on) for each file from the target, only when you browse the containing folder on your source. For example, if on the target there are two files ("one.xml" and "two.xml") and a folder "Target folder" containing a file "three.xml", should you choose to import them on demand, once you browse the root of your source, Tiger Bridge will create stub files "one.xml" and "two.xml" and "Target folder". "Target folder" will remain empty until you open it on your source and Tiger Bridge will then create a stub file "three.xml".
Import all metadata	Your source will be populated with all folders and files from the target, but the files will be represented by stub files (nearline or offline files, depending on the tier of the target they are stored on), which you can retrieve either on demand (by attempting to open them) or manually, through Tiger Bridge.
Restore all data	Tiger Bridge retrieves on your source all files and folders from the target, keeping their hierarchical structure. All retrieved files will be with replicated status. Important: The operation may take time and the free space on your source must be enough to accommodate all data found only on the target.
Link existing data	All files with matching metadata (name and size) found both on your source and the target are linked and represented as replicated. Any file for which no match is found on your source is imported as a stub file (a nearline or offline file depending on the tier of the target it is stored on).

5. Resume the work of Tiger Bridge


For the disaster recovery scenario, you can select **Import on demand**, **Import all metadata**, or **Restore all data**.

The contents of the damaged source folder will appear again when you need them and you will be able to access all your needed files. They have retained the original folder structure.

If you select **Import on demand** or **Import all metadata**, blue tags will be visible, as the files are not downloaded locally yet but still available to use. You can work with your files normally without having to download all of them upfront.

If you open a blue-tag file and then close it, the color will switch to orange as opening the file initiates its **downloading** process.

Alternatively, to create a local copy of your data again, select your files, right-click them, and choose **Retrieve data**.



They will now be marked with an orange tag, indicating successful local replication.

If you select **Restore all data**, the files will have orange tags as they will be downloaded locally.

For more information, check out our Disaster Recovery playlist [here](#).

Send Us Your Feedback

Did you find the information you were looking for?
Email us your feedback at info@tiger-technology.com.

About Tiger Technology

Founded in 2003 and headquartered in Sofia, Bulgaria and Alpharetta, GA, USA, Tiger Technology specializes in the underlying technology of hybrid cloud workflows as well as data management software solutions designed to help customers of any size, scale and industry optimize their on-premises storage and enhance their workflows through cloud services. Tiger Technology solutions are admin-friendly, non-disruptive, transparent, and highly cost-effective.

Tiger Technology brings over 15 years of expertise in developing high-performance storage solutions for the most demanding workflows, which includes cross-platform NAS/SAN file system sharing, storage, user, project, and media management.

Throughout the years and the multitude of changes in IT and the digital landscape, Tiger Technology has chosen the path of growth and continual improvement, but it has never lost sight of what it all starts with - data. The company's current focus is bridging the gap between the edge where data is generated and the cloud where data is stored and processed.

Find out more about Tiger Technology [here](#). Catch up on [our blog here](#).

Follow us on social media:

[LinkedIn](#) | [Twitter](#) | [YouTube](#) | [Facebook](#)